Available online at www.sciencedirect.com
SCIENCE DIRECT
Physica A 368 (2006) 287–293
PHYSICA A
www.elsevier.com/locate/physa

Sensitivity of the resilience of congested random networks to rolloff and offset in truncated power-law degree distributions

Randall A. LaViolette*, W.E. Beyeler, R.J. Glass, K.L. Stamber, Hamilton Link[1]

Sandia National Laboratories, P.O. Box 5800, Albuquerque NM 87185, USA

Received 23 September 2005; received in revised form 13 December 2005
Available online 23 January 2006

Abstract

Random networks were generated with the random configuration model with prescribed truncated power-law degree distributions, parameterized by an exponent, an offset, and an exponential rolloff. As a model of an attack, each network had exactly one of its highest degree nodes removed, with the result that in some cases, one or more remaining nodes became congested with the reassignment of the load. The congested nodes were then removed, and the "cascade failure" process continued until all nodes were uncongested. The ratio of the number of nodes of the largest remaining cluster to the number of nodes in the original network was taken to be a measure of the network's resiliency to highest-degree node removal. We found that the resiliency is sensitive to both rolloff and offset (but not to cutoff) in the degree distribution, and that rolloff tends to decrease resiliency while offset tends to increase it.

PACS: 89.75.−k; 89.75.Da; 89.75.Fb; 89.75.Hc

Keywords: Networks; Congestion; Scale-free; Betweenness

1. Introduction

For a wide variety of systems, the density f of the variable x of interest resembles a scale-free distribution [1], i.e.,

$$f(x) \propto x^{-\beta}. \tag{1}$$

If, for the extremes of small and large x, there were also an offset and an exponential rolloff, respectively, then a better description of data would be given by

$$f(x) \propto \exp(-\rho x)(\phi + x)^{-\beta}, \tag{2}$$

where ρ is the rolloff strength, and ϕ is the offset. With no rolloff but finite offset, f still would be asymptotically scale-free. With a finite rolloff, f is sometimes called a truncated power-law distribution

*Corresponding author. Tel.: +1 505 284 1325.
E-mail address: ralavio@sandia.gov (R.A. LaViolette).
[1]Also in Department of Computer Science, University of New Mexico, Albuquerque, NM 87131, USA.

0378-4371/$ - see front matter
doi:10.1016/j.physa.2005.12.049

(regardless of offset), e.g. Ref. [2], and describes a wide variety of systems and models, including self-avoiding random walks [3], and, of interest here, the degree distribution of finite random graphs [4,5]. A natural point of view is that ρ measures the distance to the critical point (that corresponds to zero rolloff) and the appearance of the scale-free behavior, e.g., see discussions of branching processes [6], self-organized criticality [7], and percolation processes on networks [8].

Turning now to abstract networks, the resiliency with respect to disconnection in scale-free networks, i.e., networks with the density of the degree k describing the scale-free degree distribution (1) has been widely studied, e.g., Refs. [1,8–16]. We are interested in particular in the resiliency to cascade failure of congested networks, and to that end we follow Motter [16], who studied cascade failure as a consequence of an intentional attack on a network. One of the tactics he considered was the removal of exactly one of the highest degree nodes in a scale-free network; subsequent nodes were removed only if they became congested after the removal of the first node. If the process continued, it would become a cascade failure. This approach to model network failure is different from earlier widely cited studies of network resilience (e.g., Refs. [12–14]), in which nodes were either randomly or sequentially removed from scale-free networks.

In the following, we examined the sensitivity of cascade failure to rolloff and offset in congested random networks generated with prescribed degree distributions of the form (2). Although our networks differed from those employed by Motter [16], our approach to the problem was essentially the same as his, but without his subsequent discussion of remediation strategies, which are beyond the scope of our work.

2. Calculations

The calculation consists of generating the networks, calculating the load, removing congested nodes, and monitoring any consequent disconnection. Undirected networks were generated via the random configuration model [1], with prescribed distribution f of the degree k of the form of (2), normalized on the interval $1 \leqslant k \leqslant K$, and specified by the configuration $\{\rho, \beta, \phi, K\}$, where ρ is the rolloff strength (0, 0.05, 0.1), β is the power (2.0, 2.2, 2.5, 2.8, 3.0), ϕ is the offset (0 or 2), and K is the cutoff of maximum degree (20, 40, 200). Fig. 1 shows the effect of rolloff and offset on f for $\beta = 2.5$.

For each configuration, we generated 1000 of these "random configuration" networks, as follows. We assigned degrees to initially unconnected nodes by randomly sampling from the degree distribution f. We collected random pairs of nodes and connected them if both had at least one unconnected "slot". We never connected the same pair twice. There were always so many one-degree nodes that we never failed to fill up the "slots" in the higher degree nodes. The resulting network consisted of many components, of which most were

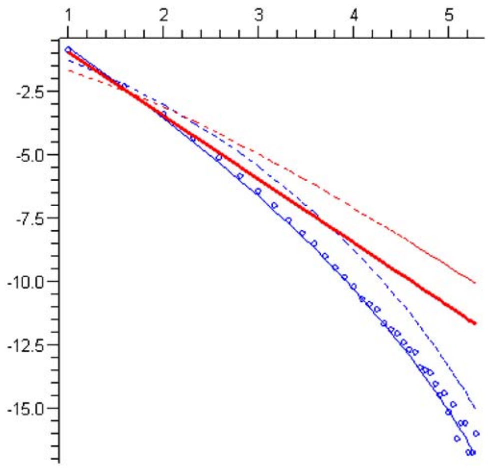

Fig. 1. The degree distributions $\log_2 f$ (vertical axis) for $K = 40$ and $\beta = 2.5$ vs. $\log_2$ degree (horizontal axis). The colors red, blue correspond to $\rho = 0$, 0.01, respectively. The line styles solid, dashed correspond to $\phi = 0$, 2, respectively. The circles correspond to the histogram of the distribution for the $\{\rho = 0.1, \phi = 2\}$ with a unit binwidth, sampled from 1000 networks with a mean number of 1000 nodes.

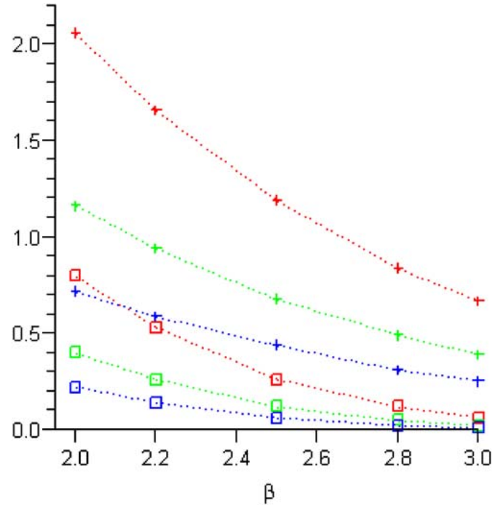

Fig. 2. Average excess edge density $\Delta\varepsilon/N$ (vertical axis) vs. β (horizontal axis). Red corresponds to zero rolloff, green and blue correspond to rolloffs $\rho = \{0.05, 0.1\}$, respectively. The open squares correspond to zero offset and the crosses correspond to offset $\phi = 2$. The standard errors are about 10%, or about twice the size of the symbols. The cutoff $K = 40$ for all the curves. The dotted lines are only a guide to the eye.

small but one was much larger than the rest, i.e., the "giant cluster", discussed in, e.g., Refs. [1,17]. Only this giant cluster became one of our trial networks. The number of nodes N of each giant cluster depends on all of the parameters in the configuration, therefore, for each configuration, we adjusted the size of the initial number of nodes so that, for each of the 1000 networks, N would fluctuate around a mean of 1000. The resulting distribution of N in the giant clusters was normal, with standard deviations in the range $10 \leqslant \sigma \leqslant 100$. Fig. 2 shows excess edge density $\Delta\varepsilon/N$ (i.e., the number of edges $-(N-1)$, or, the cycle rank of a connected graph [18], divided by N), which is zero for trees, and which decreases monotonically as β increases, corresponding to a decreasing number of cycles and an increasing tree-like structure with increasing β, for each choice of rolloff and offset. We note that the average degree (which is related to $\Delta\varepsilon$) decreases in the same way as $\Delta\varepsilon/N$. All of the networks that we generated in this way were dissortative, i.e., with a negative mixing assortativity statistic [19], and which also tracked β monotonically.

As in Motter [16], we calculated the betweenness-centrality [1,20–22] of each node (i.e., the number of shortest paths passing through that node) and called that its load. We fixed the capacity of each node to be 130% of its initial load so that a node "failed" and was removed along with its edges, if it ever became congested, i.e., subjected to a load above its capacity. Then we perturbed each network as follows: from each network we removed exactly one of its highest degree nodes, after which the load was recalculated for all of the remaining nodes. It often happened that after the removal of the first node that at least one of the remaining nodes became congested. The congested nodes were removed, and the process continued as the loads were recalculated until every remaining node was uncongested. The removal of congested nodes usually resulted in the network becoming disconnected even if only to a small extent. Also as in Motter [16], we measured the extent of the disconnection with G, the number of nodes in the largest remaining cluster divided by N, for each network. High G values (near 1) correspond to resilient networks, while lower G values correspond to more fragile networks. In many applications, which we do not pursue here, a G of much less than 0.9 would be considered to be a serious degradation of the network. As expected, G was uncorrelated to the variations in the size of the giant cluster.

3. Results

Fig. 3 shows that the distribution of G is broad and highly skewed towards low G for all configurations with finite rolloff, and even for some with no rolloff, and develops multiple modes for the less resilient cases (finite rolloff, no offset, higher β). The distribution of G is consistently narrow only for the configurations with no

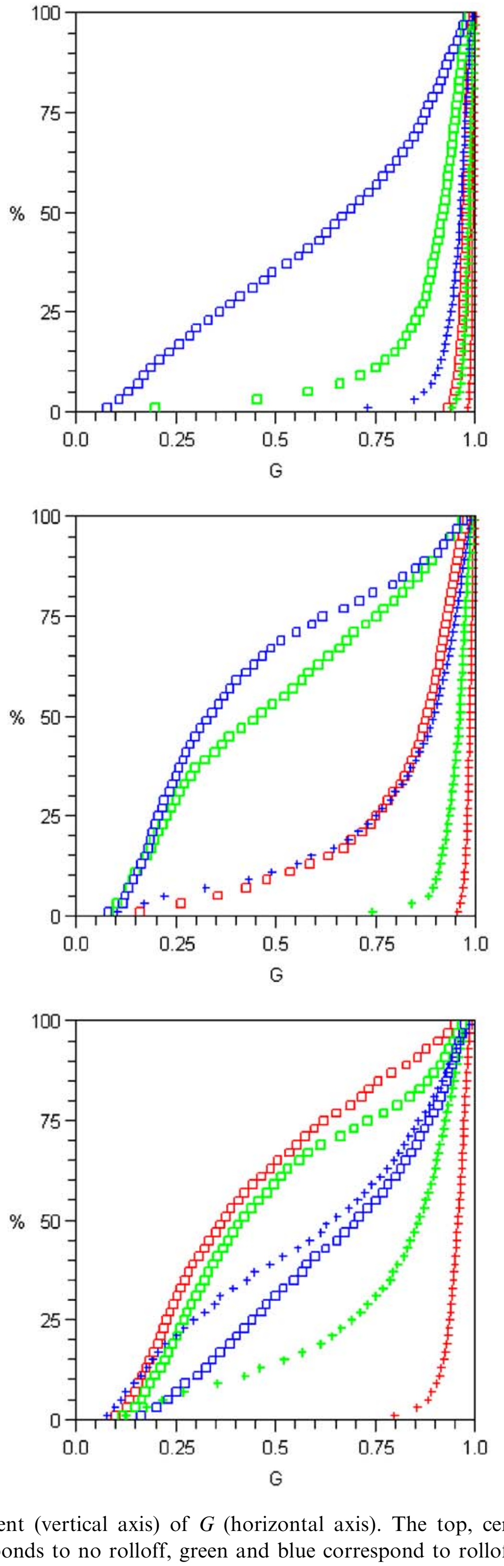

Fig. 3. Cumulative distribution in percent (vertical axis) of G (horizontal axis). The top, center, and bottom panels correspond to exponents $\beta = (2.0, 2.5, 3.0)$. Red corresponds to no rolloff, green and blue correspond to rolloffs $\rho = (0.05, 01)$, respectively. The open squares correspond to zero offset and the crosses correspond to offset $\phi = 2$. The cutoff $K = 40$ for all the curves.

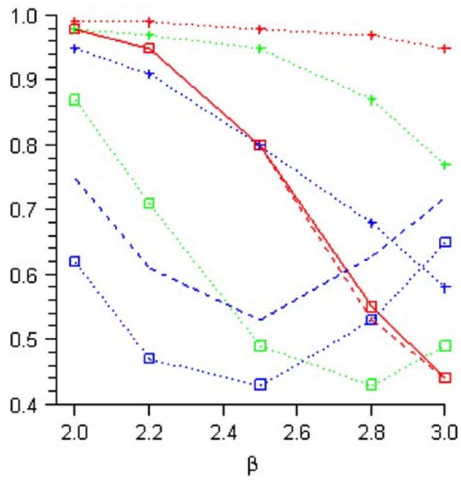

Fig. 4. Mean G (vertical axis) vs. exponent β (horizontal axis). The lines are guides to the eye; the case of no rolloff, no offset is given a solid line to aid the comparison. As in Fig. 3, red corresponds to no rolloff, green and blue correspond to rolloffs $\rho = (0.05, 0.1)$, respectively. The open squares correspond to no offset and the crosses correspond to offset $\rho = 2$. Dashed lines without symbols correspond to the configurations $\{\rho = 0, \phi = 0, K = 200\}$ and $\{\rho = 0.1, \phi = 0, K = 20\}$, respectively; the symbols correspond to $K = 40$. The standard error of the mean of G is less than 0.01.

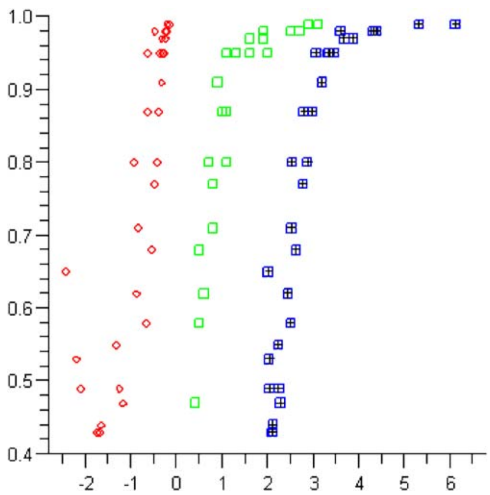

Fig. 5. The mean G (vertical axis) as a function of, respectively, the mixing assortativity ($\times 10$, with red circles), the local clustering coefficient ($\times 100$, with green squares), and the average degree (with blue squares and crosses), all read on the horizontal axis. The cutoff $K = 40$ for all of the results presented in this figure. The standard errors are 1–3% for the mixing assortativity, 10–30% for the local clustering coefficient, and no more than 0.2% for the average degree.

rolloff and finite offset. We show the mean G in Fig. 4, but the nature of the distributions make the percentiles at least as useful (if also more uncertain) a statistic [23]. The upper bound of the standard errors of the mean is 0.01. The standard error is a symmetric statistic, but more realistic uncertainty intervals for mean would be skewed towards low G values; indeed, for most configurations, we found at least one network with G below 0.2, many standard errors below the lowest mean G.

Both Figs. 3 and 4 show that the introduction of the rolloff usually decreases the resiliency of the network to disconnection after high-degree node attack, as seen in either the reduced mean G, except for ($\beta = 2.8$, 3.0; zero offset) where this behavior is unexpectedly reversed. The introduction of the finite offset on the other hand increases its resiliency (except for $\rho = 0.1$, $\beta = 3.0$). The results are not sensitive to the choice of the cutoff, at least in the following sense: making the cutoff larger does not make most resilient results much worse, making the cutoff smaller does not make the least resilient results much better, and it does not alter the qualitative variation with of G with topology.

In an attempt to begin to understand the topological origins of these results, we display in Fig. 5 the mean G vs. the following statistics for the unperturbed networks: the mixing assortativity [1,19], the local clustering coefficients [1], and the average degree, (the latter of which is also related to the excess edge density $\Delta\varepsilon/N$ by $2(\Delta\varepsilon/N+1)$, to within $1/N$). Almost any mean G between 0.4 and 0.95 (effectively almost any G for a given network) can be found for the unperturbed networks characterized by these statistics below a threshold of about 0 for the mixing assortativity, 0.01 for the local clustering coefficient, and 3 for the average degree. All three of the topological statistics studied here vary monotonically with β for fixed ρ and ϕ or with ρ for fixed β and ϕ; in particular increasing the rolloff always lowers the mixing assortativity, the average degree or edge density, and the local clustering coefficient, while increasing the offset raises all of these. On the other hand, it is apparent in Fig. 5 that the mean G achieves a minimum as a function of the mixing assortativity. Minima in the mean G vs. configuration parameters are also apparent in Fig. 4.

4. Summary and conclusions

We employed G for the figure of merit for resiliency. The introduction of the rolloff ρ to the scale-free degree distribution makes the resulting random networks less resilient to congestion-driven cascade failure (but not for $\beta = 3$; see Figs. 3 and 4), which is reflected in broader distributions of G and the lower mean G. On the other hand, the introduction of the offset ϕ makes the networks more resilient than the corresponding networks with no offset (but not for $\rho = 0.1$; $\beta = 3$; see Figs. 3 and 4). The most resilient case considered here was that of no rolloff with finite offset, for which cascades hardly developed at all. A rough explanation for these trends is that the rolloff makes the networks more sparse and tree-like, while the offset makes networks with more edges and cycles (see Fig. 2). Although those observations are correct, they do not provide a complete explanation for the trends because, as Fig. 4 shows, the mean G achieves a minimum in networks which are less tree-like, i.e., with more edges and cycles, than the most tree-like considered here ($\beta = 3$, no offset. We do not have an explanation for the appearance of the minima in Fig. 4, i.e., for why the most tree-like cases should be more resilient than those networks of moderate connectivity. Nevertheless, this behavior reminds us of Braess' paradox [24], which describes cases in which removing edges can unexpectedly improve congestion.

There appears to be a threshold in topological statistics (e.g., mixing assortativity, cluster coefficient, average degree, and excess edge density) for resilient behavior, below which a wide variety of behaviors are nearly equally probable (Fig. 5). In the absence of a theory we refrain for now from calling these empirical thresholds "critical" but we hope in future work that we might be able to identify these thresholds as such. Although cascades can, and often do, occur in any network following highest-degree node removal, these thresholds in mark a transition in their effectiveness so far as disconnection is concerned.

In conclusion, the resiliency of networks to this kind of failure was not found to be especially sensitive to the choice of maximum degree. On the other hand, the resiliency was found to be sensitive to those deviations from scale-free behavior that can be expressed with a finite offset, rolloff, or both.

Acknowledgments

The National Infrastructure Simulation and Analysis Center (NISAC) supported this work. NISAC's core partners are Sandia National Laboratories and Los Alamos National Laboratory. Sandia is a multi-program laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under contract DE-AC04-94AL85000. Los Alamos National Laboratory is operated by the University of California for the United States Department of Energy under contract W-7405 ENG-36. RAL thanks Sandia colleagues Bruce Hendrickson and Cynthia Phillips, and Los Alamos colleagues Eli Ben-Naim and Adilson Motter, for helpful discussions.

References

[1] M.E.J. Newman, SIAM Rev. 45 (2003) 167.
[2] S. Mossa, M. Barthelemy, H.E. Stanley, L.A.N. Amaral, Phys. Rev. Lett. 88 (2002) 1387011.
[3] P.-G. DeGennes, Scaling Concepts in Polymer Physics, Cornell University Press, Ithaca, 1979.
[4] M.E.J. Newman, S.H. Strogatz, D.J. Watts, Phys. Rev. E 64 (2001) 026118.
[5] R. Otter, Ann. Math. 49 (1948) 583.
[6] T.E. Harris, The Theory of Branching Processes, Dover, New York, 1989.
[7] H.J. Jensen, Self-Organized Criticality: Emergent Complex Behavior in Physical and Biological Systems, Cambridge University Press, New York, 1998.
[8] R. Cohen, D. ben-Avraham, S. Havlin, Phys. Rev. E 66 (2002) 36113.
[9] R. Albert, A.L. Barabasi, Rev. Mod. Phys. 74 (2002) 47.
[10] S. Bornholdt, H.G. Schuster (Eds.), Handbook of Graphs and Networks, Wiley, Weinheim, 2003.
[11] E. Ben-Naim, H. Frauenfelder, Z. Toroczkai (Eds.), Complex Networks, Lecture Notes in Physics, vol. 650, Springer, Berlin, 2004.
[12] R. Albert, H. Jeong, A.-L. Barabasi, Nature 406 (2000) 378.
[13] A. Broder, R. Kumar, F. Maghoul, P. Raghavan, S. Rajagopalan, R. Stata, A. Tomkins, J. Wiener, Comput. Networks 33 (2000) 309.
[14] D.S. Callaway, M.E.J. Newman, S.H. Strogatz, D.J. Watts, Phys. Rev. Lett. 85 (2000) 5468.
[15] R. Cohen, S. Havlin, D. ben-Avraham, in: S. Bornoldt, H.G. Schuster (Eds.), Handbook of Graphs and Networks, Wiley, Weinheim, 2003, p. 85.
[16] A.E. Motter, Phys. Rev. Lett. 93 (2004) 098701.
[17] M. Molloy, B. Reed, Combinat. Probab. Comput. 7 (1998) 295.
[18] F. Harary, Graph Theory, Addison-Wesley, Reading, MA, 1969.
[19] M.E.J. Newman, Phys. Rev. Lett. 89 (2002) 2087011.
[20] M.E.J. Newman, Phys. Rev. E 64 (2001) 016132/1.
[21] P. Holme, Adv. Complex Syst. 6 (2003) 163.
[22] K.-I. Goh, E. Oh, C.-M. Ghim, B. Kahng, D. Kim, in: E. Ben-Naim, H. Frauenfelder, Z. Toroczkai (Eds.), Complex Networks, vol. 650, Springer, Berlin, 2004, p. 105.
[23] A. Stuart, J.K. Ord, Kendall's Advanced Theory of Statistics, Halsted Press, New York, 1994.
[24] C.M. Penchina, L.J. Penchina, Am. J. Phys. 71 (2003) 479.